Mon. Not. R. Astron. Soc. **000**, 000–000 (0000) Printed 5 February 2008 (MN LATEX style file v1.4)

A graphical selection method for parametric models in noisy inhomogeneous regression

Nicolai Bissantz and Axel Munk

Institut für Mathematische Stochastik, Universität Göttingen, Lotzestr. 13, 37083 Göttingen, Germany

arXiv:astro-ph/0205536v1 30 May 2002

5 February 2008

ABSTRACT

A common problem in physics is to fit regression data by a parametric class of functions, and to decide whether a certain functional form allows for a good fit of the data. Common goodness of fit methods are based on the calculation of the distribution of certain statistical quantities under the assumption that the model under consideration *holds true*. This proceeding bears methodological flaws, e.g. a good "fit" - albeit the model is wrong - might be due to over-fitting, or to the fact that the chosen statistical criterion is not powerful enough against the present particular deviation between model and true regression function. This causes particular difficulties when models with different numbers of parameters are to be compared. Therefore the number of parameters is often penalised additionally. We provide a methodology which circumvents these problems to some extent. It is based on the consideration of the error distribution of the goodness of fit criterion under a broad range of possible models - and not only under the assumption that a given model holds true. We present a graphical method to decide for the most evident model from a range of parametric models of the data. The method allows to quantify statistical evidence *for* the model (up to some distance between model and true regression function) and not only *absence of evidence* against, as common goodness of fit methods do. Finally we apply our method to the problem of recovering the luminosity density of the Milky Way from a de-reddened *COBE/DIRBE* L-band map. We present statistical evidence for flaring of the stellar disc inside the solar circle.

Key words: methods: data analysis - methods: statistical - Galaxy: disc - Galaxy: structure.

1 INTRODUCTION

Often one is confronted with the problem to reconstruct an unknown function $f(t_i)$ from noisy observations $y_i = y(t_i), i = 1, \dots, N$. Astrophysical examples include reverberation mapping of gas in active galactic nuclei and recovery of the spatial (three-dimensional) luminosity density of a galaxy from blurred observations of its surface brightness. See e.g. Lucy (1994) for more examples of astronomical inverse problems. In this paper we are concerned with a new method to compare several competing parametric models for the regression function f.

Due to the noisy measurements it is tempting to assume that $y_i = f(t_i) + \varepsilon_i$, where the ε_i denote some random noise and $f(t_i)$ the expected value of y_i, i.e. $E[y_i] = f(t_i)$. In particular we allow for different error distributions of the ε_i, which entails inhomogeneous variance patterns, viz. $V[\varepsilon_i] = \sigma_i^2$, as will be the case in our example of de-projecting the de-reddened *COBE/DIRBE* L-band surface brightness map of Spergel et al. (1996), as discussed by Bissantz & Munk (2001, [BM1]).

It is a common proceeding to fit a class of functions $U = \{f_\vartheta : \vartheta \in \Theta\}$ (parametric model) to the data y_i. The parametric model may depend on a parameter ϑ, where $\vartheta \in \Theta \subseteq \mathbb{R}^d$. A popular method to select a "best-fitting" ϑ from Θ is to minimise the empirical mean squared error (MSE)

$$Q_N^2(\vartheta) := \sum_{i=1}^{N} (y_i - f_\vartheta(t_i))^2 \qquad (1)$$

or weighted variants of it. This gives $\hat{\vartheta}$, the least squares

© 0000 RAS

estimator (LSE) of $\hat{\vartheta}$. Other measures of goodness of fit are e.g. L^1-error criteria, where the absolute deviation between y_i and $f_\vartheta(t_i)$ is considered (Seber & Wild 1989). A small value of $Q_N^2(\hat{\vartheta})$ often is used as an indication for a good explanation of the observations by the model $f_{\hat{\vartheta}}$. Note that in regression models where the noise is *inhomogeneous* the quantity $Q_N^2(\hat{\vartheta})$ is often useless (cf. [BM1] for an explanation) and more subtle methods have to be applied.

An advantage of the parametric fitting methodology in contrast to nonparametric curve estimation, i.e. approximating the data by arbitrary functions (e.g. by splines, orthogonal series or wavelets, cf. Efromovich, 1999 or Hart, 1997) is the fact that often physical reasoning resulting from a theory suggests such a class of functions U. Furthermore, subsequent data analysis and interpretation becomes very simple if once a proper $f_{\hat{\vartheta}}$ is selected. Hence it is an important task to pre-specify U correctly in order to obtain a reasonable fit. Therefore in this paper we discuss the problem of evaluating the goodness of fit of a parametric model U. Moreover, we offer a graphical method which allows to select a proper model U from a class of different models $\mathcal{U} = \{U_k\}_{k=1,\dots,l}$, say.

A common proceeding is to assume that the model holds, and to test if the observed data give reason to reject the model. This type of goodness of fit tests is performed by evaluating the probability distribution of a pre-specified measure of discrepancy, such as $Q_N^2(\hat{\vartheta})$. This is done under the assumption that U holds true. Then, when this measure exceeds a certain quantity, the model U is rejected.

One problem of such methods is that a large data set leads essentially to rejection of any model U (an illustrative discussion can be found in Berger, 1985), because the "real world" is never exactly described by such a model and as the number of observations increases, statistical methods will always detect these deviations between the model and "reality". Conversely, the selected statistical criterion may lead to a decision in favour of U (albeit wrong), because it is not capable to detect important deviations from U or the decision is affected by quantities which are not captured in the model U (e.g. correlation between the y_i). Another problem can be over-fitting of data by models with a too large number of parameters. Therefore, various methods have been suggested which penalise the number of parameters, i.e. the complexity of a model (Akaike, 1974, Burnham et al., 1998, or Schwarz, 1978).

In this paper, we suggest a methodology which aims to avoid these problems by considering the distribution of a discrepancy measure such as $Q_N^2(\hat{\vartheta})$ under all "possible" functions f. This extends the method given in [BM1] to the more realistic case where the "true" function f is not restricted to be in U. Furthermore, a graphical method will be presented which allows to select the most appropriate between several competing models U_i. With our method, this is still possible if these models have different numbers of parameters.

In the next section we will describe the method and its algorithmic implementation, the wild bootstrap. Based on the theory presented in sect. 2, we suggest in sect. 3 a graphical method to assess the validity of U as well as to compare between different models. This method is denoted as p-value curve analysis. In sect. 4 our method is applied to a near-infrared [NIR] L-band map of the Milky Way [MW] and two different models of the spatial luminosity distribution are compared. One of the models includes a flaring disc component. We analyse the models' p-value curves, and find that flaring in the disc improves the fit to the data.

2 A NEW METHOD OF MODEL SELECTION

In section 2.1 we briefly recall the methodology suggested in [BM1] and extend it to the situation where f is not in the model U. This will be used to compute p-value curves, a graphical method of model diagnostics, which was introduced by Munk & Czado (1998) in a different context. In sect. 2.2 we describe the practical application of the method.

2.1 Basic theory of the method

We begin with an introduction to the basic principles of our method. As mentioned above $Q_N^2(\hat{\vartheta})$ fails to be a valid criterion for goodness of fit in inhomogeneous models [BM1]. Instead we replace the pure residuals $y_i - f_{\hat{\vartheta}}(t_i)$ with smoothed residuals, to allow for a valid statistical analysis. For the smoothing step we require an *injective* linear integral operator with kernel T, viz.

$$g(w) = \mathbf{T}(f)(w) = \int T(w,v) f(v) dv$$

which maps the function f to be recovered onto g. In principle any injective operator $\mathbf{T}$ is a valid option for the smoothing, however a good choice is driven by aspects such as efficiency and simplicity. In our example (cf. sect. 4) we introduce "cumulative smoothing" with $T(w,v) = \min(w,v)$. An extensive simulation study by Munk & Ruymgaart (1999) revealed this smoothing kernel as a reasonable choice which yields a procedure capable to detect a broad range of deviations from the class of functions $U = \{f_\vartheta : \vartheta \in \Theta\}$.

A measure of the discrepancy between the "true" f and U is the transformed distance

$$D^2(f) = \min_{\vartheta \in \Theta} ||\mathbf{T}(f - f_\vartheta)||^2 \tag{2}$$

where the norm refers to some L^2-norm. Now assume that the minimum in eq. 2 is achieved at a parameter vector $\vartheta^* = \vartheta^*(g) \in \Theta$. Because ϑ^* is unknown it has to be estimated from the data. This can be done by numerical minimisation of the empirical counterpart of the r.h.s. of eq. 2,

$$\hat{D}^2 := \min_{\vartheta \in \Theta} ||\mathbf{T} f_\vartheta - \hat{g}||^2$$

where

$$\hat{g} = N^{-1} \sum_{i=1}^{N} y_i T(u, t_i)$$

is an estimation of g using the noisy data y_i. The reasoning behind this approach is that for sufficiently large number N of observations (in our example $N = 4800$) it can be shown that $\hat{g}$ converges in probability to the true (but unknown) function g, independently whether a parametric model U is valid or not. On the other hand the empirical minimiser $\hat{\vartheta}_T$ estimates the best possible fit to $\hat{g}$ by the model U. The resulting estimator is denoted as a smoothed minimum distance estimator $\hat{\vartheta}_{\mathbf{T}}$ (SMDE) and has the property that, if the true function $f = g_{\vartheta^*}$ is in U, $\hat{\vartheta}_{\mathbf{T}} \to \vartheta^*$ as the sample size increases. For detailed proofs we refer to Munk & Ruymgaart (1999).

Note that ϑ^* is the "true" best-fitting parameter vector, which could only be determined if the data would be free of noise, whereas $\hat{\vartheta}_{\mathbf{T}}$ is an estimation of the best-fitting parameter vector using the noisy data. Here and in the following, quantities with a hat, " ˆ ", are estimated from the noisy data, whereas such without a hat are the "true" functions to be recovered.

Munk & Ruymgaart (1999) showed that the probabilistic limiting behaviour of $\hat{D}^2$ depends on whether f belongs to the model U under investigation. More precisely when f belongs to U the distribution of $N\hat{D}^2$ is for large N approximately that of

$$\sum_{i=1}^{\infty} \lambda_i \chi_i^2 \tag{3}$$

where χ_i^2 denotes a sequence of independent squares of standard normal random variables and $\lambda_i \geq 0$ is a sequence of real numbers, s.t. $\sum_{i=1}^{\infty} \lambda_i^2 < \infty$, which depend on ϑ^*, the distribution $\mathcal{L}$ of errors ε_i and the operator T.

In contrast if f does not belong to U, we have $D^2 > 0$ and $N^{\frac{1}{2}} \left(\hat{D}^2 - D^2\right)$ tends for large N to a centred normal distribution with variance $\sigma^2_{\mathbf{T},\mathcal{L},\vartheta^*}$, depending on $\mathbf{T}$, ϑ^*, and $\mathcal{L}$. Observe, that we obtain two different types of distributions, accordingly to the situation whether the "true" (unknown) function f is in the model U or not. Because of the complicated dependency of the $(\lambda_i)_{i \in I\!N}$ and $\sigma^2_{\mathbf{T},\mathcal{L},\vartheta^*}$ on $\mathbf{T}$, ϑ^*, and $\mathcal{L}$ a resampling algorithm should be applied in order to approximate these limiting distributions. Stute et al. (1998) presented a wild bootstrap algorithm which can be used to approximate the law $N\hat{D}^2$. Munk (1999) showed that this algorithm is also valid when f does not belong to U, which is crucial for our paper. This algorithm will be carefully explained in the next paragraph. Recall that the subsequent bootstrap algorithm allows to determine the probability distribution of the quantity of interest $\hat{D}^2$. The general strategy of our method will be the following. Because $\hat{D}^2$ measures the distance between the model U and the estimator $\hat{g}$ from noisy data, knowledge of the probability distribution of $\hat{D}^2$ (which will be determined by the subsequent bootstrap algorithm) allows us to quantify whether an observed value of $\hat{D}^2$ for a model U is more likely than for a competing model U', say. Even, when none of these models is completely true (which is always the case in the real world) $\hat{D}^2$ quantifies the best possible *approximation* of g by U or U' respectively.

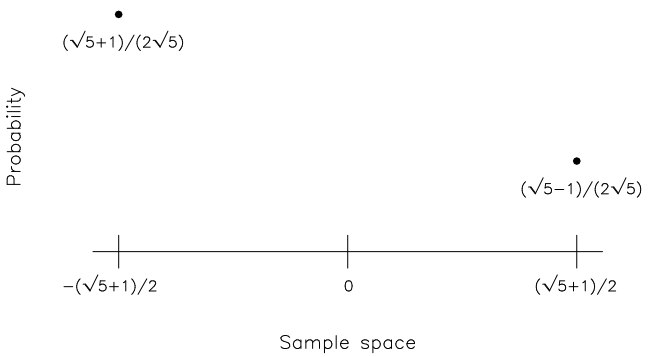

Figure 1. Binary probability distribution required in step 2 of the wild bootstrap algorithm. The ordinate gives the probability of the random number to be $-(\sqrt{5}+1)/2$ and $(\sqrt{5}+1)/2$, respectively.

2.2 Practical application of the method

We now introduce the resampling algorithm to approximate the law $N\hat{D}^2$. The algorithm starts with the determination of the SMDE $\hat{\vartheta}_{\mathbf{T}}$ and the smoothed residuals between this model and the data (step 1). Then in step 2-5 the resampling part of the algorithm follows. The same algorithm is used in [BM1].

Step 1: (*Generate residuals*). Compute residuals

$$\hat{\varepsilon}_i := y_i - f_{\hat{\vartheta}_{\mathbf{T}}}(t_i), \qquad i = 1, \cdots, n$$

where $\hat{\vartheta}_{\mathbf{T}}$ denotes a solution of the minimisation of

$$\hat{D}^2 := \chi^2(\hat{\vartheta}_{\mathbf{T}}) := \min_{\vartheta \in \Theta} \|\hat{g} - \mathbf{T} f_\vartheta\|^2.$$

Step 2: (*The "wild" part*). Generate new random variables c_i^*, $i = 1, \ldots, n$, which do *not* depend on the data, where each c_i^* is distributed to a distribution which assigns probability $(\sqrt{5}+1)/2\sqrt{5}$ to the value $(-\sqrt{5}-1)/2$ and $(\sqrt{5}-1)/2\sqrt{5}$ to the value $(\sqrt{5}+1)/2$. See fig. 1 for a visualisation of this probability distribution.

Step 3: (*Bootstrapping residuals*). Compute $\varepsilon_i^* := \hat{\varepsilon}_i c_i^*$ and $y_i^* = f_{\hat{\vartheta}_{\mathbf{T}}} + \varepsilon_i^*$. This gives a new data vector $(y_i^*, t_i)_{i=1,\ldots,n}$.

Step 4: (*Compute the target*). Compute $\hat{D}_*^2$ with $(y_i^*, t_i)_{i=1,\ldots,n}$

Step 5: (*Bootstrap replication*). Repeat step 1-4 B times which gives values $\hat{D}_{1,*}^2, \ldots, \hat{D}_{B,*}^2$. B is a large number, typically $B = 500$ or $B = 1000$ is sufficient.

From the bootstrap replications $\hat{D}^2_{1,*},\dots,\hat{D}^2_{B,*}$ we compute the quantities

$$x_1 = \sqrt{N}\left(\hat{D}^2_{1,*} - \hat{D}^2\right),\dots,x_B = \sqrt{N}\left(\hat{D}^2_{B,*} - \hat{D}^2\right),$$

using the number of data points N. The $x_1,\dots,x_B$ are realisations of the random quantity $X = \sqrt{N}\left(\hat{D}^2_* - \hat{D}^2\right)$. It can be proved that the empirical distribution function of $\hat{D}^2_{1,*},\dots,\hat{D}^2_{B,*}$ yields an approximation to the true distribution of $\hat{D}^2$ after a proper re-centring, i.e. the cumulative probability distribution function F^*_B of $X = \sqrt{N}\left(\hat{D}^2_* - \hat{D}^2\right)$ is close to the cumulative distribution function of$\sqrt{N}\left(\hat{D}^2 - D^2\right)$ for any $D^2 > 0$ (Munk, 1999).

An important application of this result is to determine an approximation to the probability $p(t, D^2)$ that $\hat{D}^2$ is below a certain value t, provided the distance between true function f and the model is D^2. To this end we use that F^*_B, found from the bootstrap replications, approximates the (unknown) cumulative probability distribution of $\sqrt{N}(\hat{D}^2 - D^2)$. The latter distribution allows to determine $p(t, D^2)$. Hence we are in the position to compare the probability that the observed value of $\hat{D}^2$ is achieved in all "possible worlds", i.e. for any possible f. In fact, it turns out that this probability does only depend on f via $D^2(f)$, which allows a nice geometric interpretation as we will illustrate in the following.

We will use the asymptotic similarity of the two cumulative probability distributions in the following section to estimate the probability $p(t, D^2)$. From this we then define the P-value curve $\alpha_N(\Pi)$, which can be regarded as a measure of evidence for $D^2 \le \Pi$, given $\hat{D}^2$ and F^*_B. Thus these quantities allow to constrain D^2 for a parametric model of a given set of data.

3 *P*-VALUE CURVES

The main methodology we propose in this paper is the computation of a p-value curve as a graphical tool for illustrating the evidence of a model. To this end we plot the function $\alpha_N(\Pi) = F^*_B\left(\sqrt{N}\left(\hat{D}^2 - \Pi\right)\right)$ for $\Pi > 0$, i.e. the value of $\alpha_N(\Pi)$ is given by the probability that the random quantity $X = \sqrt{N}\left(\hat{D}^2_* - \hat{D}^2\right)$ is smaller than $\sqrt{N}\left(\hat{D}^2 - \Pi\right)$. Note that this implies that for Π increasing $\alpha_N(\Pi)$ decreases, because we then evaluate the cumulative distribution function $F^*_B(x)$ for decreasing x, and in particular, if $\alpha_N(\Pi)$ is small, at the left tail of F^*_B.

The interpretation of the function $\alpha_N(\Pi)$ is as follows. Assume the true distance between model U and f (i.e. the distance between the minimising f_{ϑ^*} and the "true" function f) is $D^2 = \Pi$. If this holds, the probability that$\sqrt{N}(\hat{D}^2 - D^2)$ is smaller than some value t is given as

$$P_{D^2=\Pi}\left(\sqrt{N}\left(\hat{D}^2 - D^2\right) \le t\right) \approx F^*_B(t) \tag{4}$$

where the r.h.s. denotes the bootstrap approximation to the true distribution function on the l.h.s. Now we reject the hypotheses $H : D^2 > \Pi$ (vs. alternative $K : D^2 \le \Pi$) whenever $\alpha_N(\Pi) \le \alpha$ for a given level of significance α. Hence $1 - \alpha_N(\Pi)$ can be regarded as the estimated evidence in favour of the model U (up to a distance between model and data $D^2 \le \Pi$).

Note that this approach highlights the fact that finally the astrophysicist has to decide whether a value of $D^2 = \Pi$ should be regarded as scientifically negligible or as deviation from the model U which is considered as too large by astrophysical reasons. We mention that the classical goodness of fit tests do not offer the scientist the specification of such a value Π.

How can an upper bound for a just acceptable D^2 be determined? One simple suggestion is to compute the distance $\tilde{D}^2 = ||\mathbf{T}f_{\hat{\vartheta}_\mathbf{T}} - \mathbf{T}\tilde{f}_{\hat{\vartheta}_\mathbf{T}}||^2$ between the best model $f_{\hat{\vartheta}_\mathbf{T}}$ and "test models" $\tilde{f}_{\hat{\vartheta}_\mathbf{T}}$. Such test models should then be constructed from $f_{\hat{\vartheta}_\mathbf{T}}$ by adding (systematic) deviations to the model, which are still considered as scientifically negligible differences to the best model. Then, if D^2 is not larger than the average over the test models $< \tilde{D}^2 >$, computed from a number of such test models, it is considered as scientifically negligible.

Observe that with our proposed method the statistical type one error is the error to decide for the model (or more precise for a neighbourhood $D^2 \le \Pi$ of the model) although it is not valid. Classical goodness of fit tests are only able to control the error of rejecting the model albeit it holds, i.e. they are based on testing $H_0 : D^2 = 0$ vs. $K_0 : D^2 > 0$.

Fixing $\hat{D}^2$, a small value of $\alpha_N(\Pi)$ indicates large probability for $D^2 \le \Pi$ and a large value (close to 1) of $\alpha_N(\Pi)$ indicates a large probability for $D^2 > \Pi$. It is important to note that the interesting regions of the resulting curves $\alpha_N(\Pi)$ are those values of Π where $\alpha_N(\Pi)$ is rather large (larger than 0.9 say) and rather small (smaller than 0.1) in accordance with the usual choice of levels of significance. In contrast decisions based on α_N in regions where $\alpha_N(\Pi) \approx 0.5$ would correspond to flipping a coin in order to decide whether $D^2 \le \Pi$ or not.

As an important advantage of p-value curves we find that it gives us not only an estimated probability (p-value) that we would observe a test statistic (such as $Q^2_N(\hat{\vartheta})$ or $\hat{D}^2$) provided the assumption that U underlies the data is true. Rather we obtain simultaneously all scenarios over the entire range of "possible worlds" which are parametrised by D^2. In particular this implies that models with a large number of parameters are penalised in an automatic way. As the number of parameters increases the variability of the statistic $\hat{D}^2$ increases and hence the variability of F^*_B, i.e. the range of values for X, for which F^*_B differs significantly from 0 and 1, is larger. On the other hand the bias is reduced. As the number of parameters decrease the opposite will be the case. This leads to a curve $\alpha_n(\Pi)$ which slowly decreases to zero if the variance is too large or if the bias is too large.

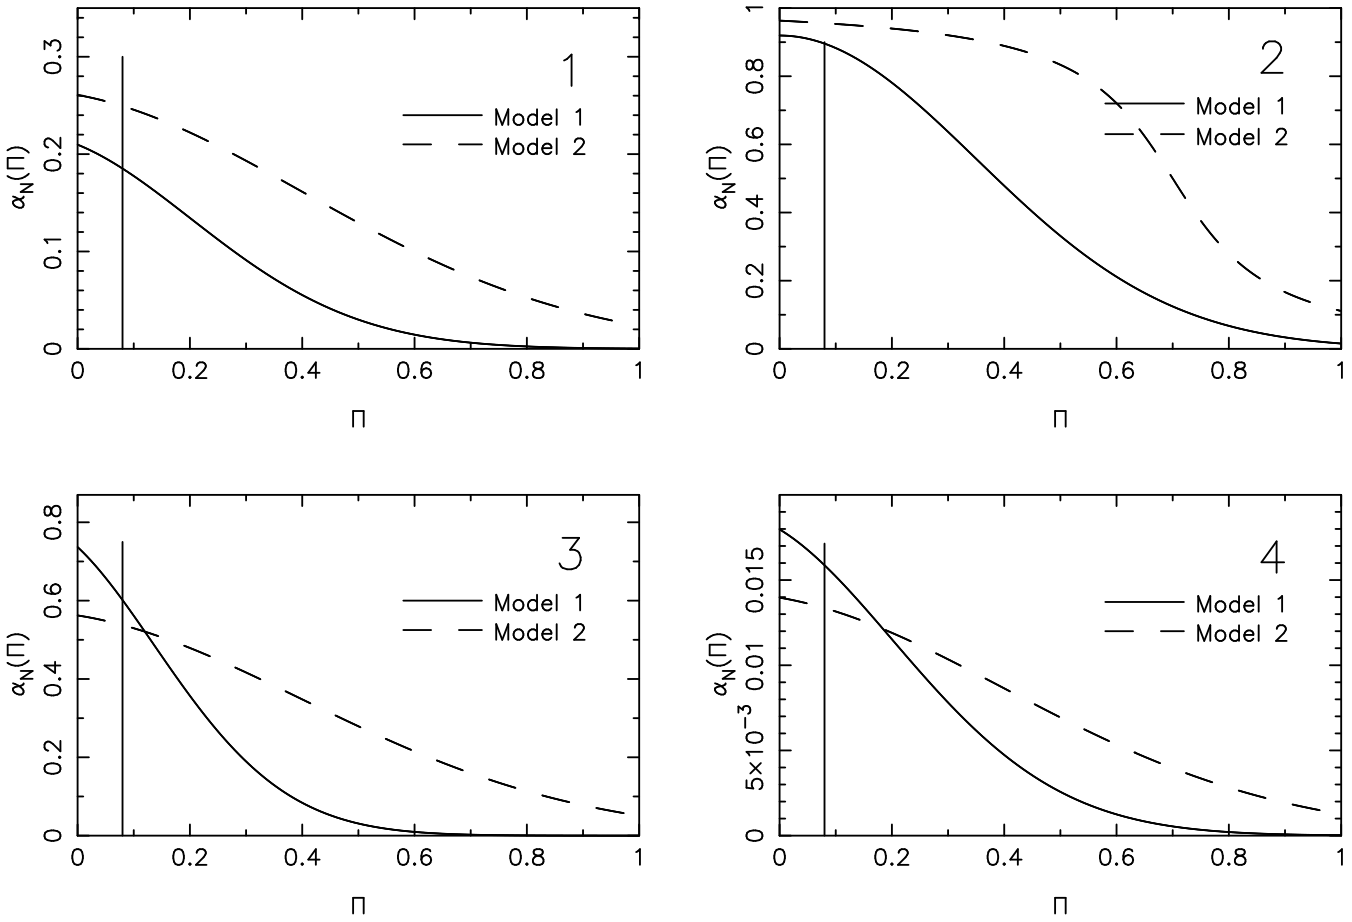

Figure 2. Typical cases for p-value curve comparison of two parametric models. The vertical lines at $\Pi = 0.08$ in the graphs indicate the observed value of $\hat{D}^2 = 0.08$. In graph 1, model 1 fits better, as well as in graph 2. However in graph 2 is additionally strong evidence that model 1 does not hold. Graph 3 is again an example with model 1 the better model. Finally in graph 4 the situation depends on the assumption of the distance between the parametric model U and the true regression function f (cf. sect. 2) D^2.

Hence evidence for a small Π can only be claimed if these two quantities are balanced.

In other words a p-value curve reflects automatically the tradeoff between variance and bias in a regression. Here the bias of the regression functions can be viewed as the difference between the "true" expectation value $E[y_i]$ and the value of the regression function $f(t_i)$. The variance provides an estimate of the uncertainty of the best-fitting parameters $\hat{\vartheta}$ or $\hat{\vartheta}_{\mathbf{T}}$.

Before we analyse two competing models for the structure of the MW we illustrate in an artifical example typical features of p-value curves. In fig. 2 various scenarios are displayed. In graph 1 model 1 beats model 2 at all fronts. The estimated evidence for $D^2 \leq \Pi$ is uniformly larger for any $\Pi > 0$. This coincides with "classical testing" because also the classical p-value for testing H:$D^2 = 0$ is larger. Observe that the classical p-value corresponds in this graph to $1 - \alpha_N(0)$.

Graph 2 is similar, observe however, that a classical analysis would indicate that here is additionally strong evidence that model 1 does not hold ($\alpha_N(0) \gtrsim 0.9$), although it yields a better fit as model 2, exactly as in graph 1. Here the value of Π where $\alpha_N(\Pi) = 0.1$ (i.e. where $\Pi \approx 0.7$), say, becomes important because it gives an idea of the order of magnitude between model U and the true regression. Hence it has to be decided for the particular problem whether a distance of $\Pi \approx 0.7$ is considered as "large" or scientifically irrelevant.

Graph 3 represents a typical case of over-fitting by model 2. Classical reasoning would prefer model 2 because $\alpha_N(0)$ is smaller and hence the classical p-value larger. However, we see that this is due to a lack of power of the used test statistic, because the slope of the curve is very flat due to a large variability of the test statistic. Hence there is not much support for the decision $D^2 \leq 0.5$, say, ($\alpha_N(0.5) \approx 0.3$) whereas model 1 yields $\alpha_N(0.5) \approx 0.03$. Thus there is strong evidence that the distance between model 1 and the true regression curve is smaller than 0.5, say.

Finally in graph 4 both models are acceptable with slight preference to model 2 provided a distance of 0.2 (the point of intersection of both curves) is considered as an acceptable distance between U and f. If a larger distance, $\Pi = 0.5$, say is considered to be tolerable, however model 1 has to be preferred.

4 FLARING OF THE STELLAR DISC

Observations have shown that the HI disc of the MW flares (see, for example, Merrifield, 1992, or Malhotra, 1995). The situation is much less clear for the stellar disc. Alard (2000) finds flaring for the disc outwards of the solar orbit, from an analysis of 2 micron sky survey (2MASS) data, with a vertical scale-height $\approx 300\,\mathrm{pc}$ in the solar neighbourhood. Other evidence comes from Kent et al. (1991), who have fitted parametric models to Spacelab2 IR telescope (IRT) $2.4\mu m$ observations of the MW. The vertical scale-height of their best model's disc is constant in the inner $\approx 5\,\mathrm{kpc}$ with $h_z = 165\,\mathrm{pc}$, but rises outside of this galactocentric radius to $h_z \approx 247\,\mathrm{pc}$ in the solar neighbourhood. Thus the results of Alard and Kent et al. for the vertical disc scale-height in the solar neighbourhood are consistent to within $\approx 20\%$.

We apply our proposed method to dust-corrected *COBE/DIRBE* L-band data (Weiland, 1994; Spergel et al., 1996), and investigate whether there is evidence for flaring of the disc inside the solar orbit. We remark that this L-band observations are expected to trace the density of stars (Binney, Gerhard & Spergel, 1997 [BGS]). Note that from non-parametric models of this L-band data [BGS] have found vertical scale-heights $z_0 \approx 120 - 150\,\mathrm{pc}$ at $R = 5\,\mathrm{kpc}$ from the galactic centre. They remark that this is inconsistent with the value of 300 pc from star counts at the Galactic poles (Gilmore & Reid, 1983), but consistent with the findings of Kent et al. (1991). We will apply our proposed statistical test on the same data to demonstrate its ability in this context as an example application.

The general outline of this sect. is as follows: First we introduce the observational data (sect. 4.1), and construct functional forms for two different parametric models of the MW luminosity density distribution (sect. 4.2). Then (sect. 4.3) we fit these models to the *COBE/DIRBE* L-band data and apply the wild bootstrap algorithm to both models, with $B = 5000$. Finally we analyse the distribution of the distances $\hat{D}^2$ between the models and the data, both under the assumption that the respective parametric model does reproduce the data, and that this is not the case (sect. 4.4).

4.1 Observational data

The *DIRBE* experiment on board the *COBE* satellite, launched in *1989*, has provided maps of the sky in several infrared wavebands (Weiland et al. (1994)). This data has been used to estimate the luminosity distribution of the MW, both parametrically (e.g. Freudenreich, 1998, and Dwek et al., 1995), and non-parametrically (Binney & Gerhard, 1996, [BGS], Bissantz et al., 1997, and Bissantz & Gerhard, 2001).

In this paper we use a *COBE/DIRBE* NIR L-band map, corrected for dust absorption by Spergel et al. ((1996)). The resolution of the equidistant grid of data is $n = 120$ points in $-89.25\,\mathrm{deg} \le l \le 89.25\,\mathrm{deg}$ and $m = 40$ points in $-29.25\,\mathrm{deg} \le b \le 29.25\,\mathrm{deg}$. We only use the data $-60\,\mathrm{deg} \le l \le 60\,\mathrm{deg}$, $-20\,\mathrm{deg} \le b \le 10\,\mathrm{deg}$, to downweight those parts of the sky where non-informative parts in the data can be observed due to extreme noise (cf. [BM1]).

This dataset is well suited to demonstrate our proposed method since it consists of several thousand data points, enough to make the method applicable. Simulations have shown that the method is already applicable when more than 50 data points are available provided the error distribution behaves well.

4.2 The parametric models

We construct two different parametric models, one including flaring, according to the approach of Kent et al. (1991), the other not. In this section the functional forms of the models are presented, first the individual bulge and disc components. We use a Cartesian coordinate system with axes x, y, z. Here x is along the major axis, and y along the minor axis of the bulge/bar, both in the main plane of the MW. We set the position of the sun in this coordinate system to a distance from the main plane of the disc $z_\odot = 14\,\mathrm{pc}$, the distance to the galactic centre $R_\odot = 8\,\mathrm{kpc}$, and the angle between the major axis of the bar and the line-of-sight from the sun to the galactic centre $\phi_{\mathrm{bar}} = 20\,\mathrm{deg}$ ([BGS]). Let $a^2 \equiv x^2 + \left(\frac{y}{\eta}\right)^2 + \left(\frac{z}{\zeta}\right)^2$ and $r^2 \equiv x^2 + y^2$. Then the model components are:

"BGS" bar/bulge: The bulge model is selected similar to [BGS]. It is a truncated power law bulge:

$$\rho\,(x,y,z) = b \cdot \frac{e^{-a^2/a_m^2}}{a_m^3 \eta\zeta\,(1 + a/a_c)^q}$$

"BGS" disc: A double-exponential disc, without flaring [BGS]:

$$\rho\,(x,y,z) = d \cdot \left(e^{-|z|/z_0}/z_0 + \alpha e^{-|z|/z_1}/z_1\right) \cdot r_d e^{-r/r_d}$$

"Kent" disc: A double exponential disc, similar to the "BGS" disc. But now we include flaring, in spirit of the flaring disc model of Kent et al. (1991). Inside of a galactocentric radius $R_i = 5\,\mathrm{kpc}$ the scale-height h_z is constant. Outside of R_i it rises linearly to the solar neighbourhood, where the scale-height is 247 pc. We also set the radial disc scale length to the Kent et al. value of $r_d = 3.001$ kpc outside R_i. Inside R_i the radial scale length is a fit variable.
Thus we define for $r \le 5\,\mathrm{kpc}$:

$$\begin{aligned}\rho\,(x,y,z) &= d \cdot \left(e^{-|z|/z_0}/z_0 + \alpha e^{-|z|/z_1}/z_1\right) \\ &\quad \cdot r_d e^{-r/r_d}\end{aligned}$$

for $r > 5\,\mathrm{kpc}$, with $\sigma \equiv z_0 + (0.247\,\mathrm{kpc} - z_0) \cdot \frac{r/[\mathrm{kpc}]-5}{3}$:

$$\begin{aligned}\rho\,(x,y,z) &= d \cdot \left(e^{-|z|/\sigma}/\sigma + \alpha e^{-|z|/z_1}/z_1\right) \\ &\quad \cdot r_d e^{-r/3.001\,\mathrm{kpc}} e^{5(3.001^{-1} - r_d^{-1}[\mathrm{kpc}])}\end{aligned}$$

We remark that this definition of the disc ensures $\rho \in$

$C^0\left(\mathbb{R}^3\right)$, in particular that ρ is continuous at $\{(x,y,z) : r = \sqrt{x^2+y^2} = 5\,\mathrm{kpc}\}$.

We remark that we assume a priori some of the parameters as fixed. These are the disc parameters $\alpha = 0.27$ and scale-height $z_1 = 42\,\mathrm{pc}$, and the cusp parameters $a_c = 0.1\,\mathrm{kpc}$ and $q = 1.8$ [BGS]. We refer to Kent et al. (1991), [BM1] and [BGS] for a more detailed description of the models and their parameters.

Using these model parts we define in table 1 two models which we analyse.

4.3 The fitting algorithm

The algorithm in order to find estimates for the parameters in the above mentioned models is discussed in detail in [BM1] and only briefly summarised here.

Our task is to fit the de-reddened *COBE/DIRBE* L-band surface brightness map $Y_{ij} = \omega^{\mathrm{obs}}(l_i, b_j)$ (Spergel et al. 1996), which is blurred by some random error ε_{ij} at position (l_i, b_j). Particularly, an explorative data analysis shows that it is necessary to allow for a position dependent noise $\mathrm{Var}[\varepsilon_{i,j}] = \sigma^2_{i,j}$ (cf. [BM1]). The linear integral operator $\mathcal{P}$ projects a three-dimensional luminosity distribution $\rho(x,y,z)$ to a surface-brightness distribution $\omega(l,b)$ at the sky, viz:

$$\omega(l,b) = \mathcal{P}(\rho)(l,b) = \int_0^\infty \tilde{\rho}(r,l,b)dr,$$

with $\tilde{\rho}(r,l,b) \equiv \rho(x(r,l,b), y(r,l,b), z(r,l,b))$ where $\tilde{\rho}$ is defined in [BM1]. We assume that $\mathcal{P}$ is injective in a neighbourhood of $U = \{\rho_\vartheta(\cdot)\}_{\vartheta\in\Theta}$. This depends on a proper selection of the parametric model U. Thus the problem to solve is to recover the MW luminosity density ρ^{MW} from the noisy integral equation $\omega^{\mathrm{obs}}(l_i,b_j) \equiv \omega^{MW}(l_i,b_j) + \varepsilon_{ij} = \mathcal{P}\left(\rho^{MW}\right)(l_i,b_j) + \varepsilon_{ij}$. Note that ω^{MW} is the noise-free surface brightness distribution of the MW.

Following the method proposed in [BM1], let $\omega_\vartheta(l,b) = \mathcal{P}(\rho_\vartheta)(l,b)$; $\vartheta\in\Theta$, and consider the transformed model

$$U_{\mathbf{T}} = \mathbf{T}U = \{\mathbf{T}\omega_\vartheta(l,b)\}_{\vartheta\in\Theta},$$

with

$$(\mathbf{T}\omega)(u,v) = \int\int \omega(l,b)T((u,v),(l,b))dldb.$$

Here $\mathbf{T}$ is a smoothing integral operator with kernel $T((u,v),(l,b)) = \min\{u,l\}\cdot\min\{v,b\}$; $(u,v),(l,b)\in\mathbb{R}^2$. Munk & Ryumgaart (1999) have shown that this smoothing kernel is a reasonable choice (cf. sect. 2 and [BM1]).

According to sect. 2 we estimate the smoothed MW surface brightness $g^{MW}(u,v) = \left(\mathbf{T}\omega^{MW}\right)(u,v)$ from the noisy observations $Y_{ij} = \omega^{\mathrm{obs}}(l_i,b_j)$ as

$$\hat{g}^{MW}(u,v) = \frac{1}{n\cdot m}\sum_{i=1}^{n}\sum_{j=1}^{m}\omega^{\mathrm{obs}}(l_i,b_j)T((u,v),(l_i,b_j)),$$

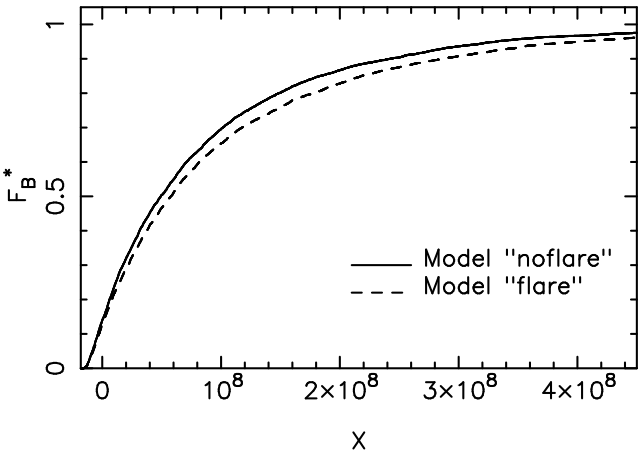

Figure 3. Distribution F^*_B of X where $X = \sqrt{N}(\hat{D}^2_* - \hat{D}^2)$ for our models of the *COBE/DIRBE* L-band data. The solid line corresponds to model "noflare", the dashed line to model "flare".

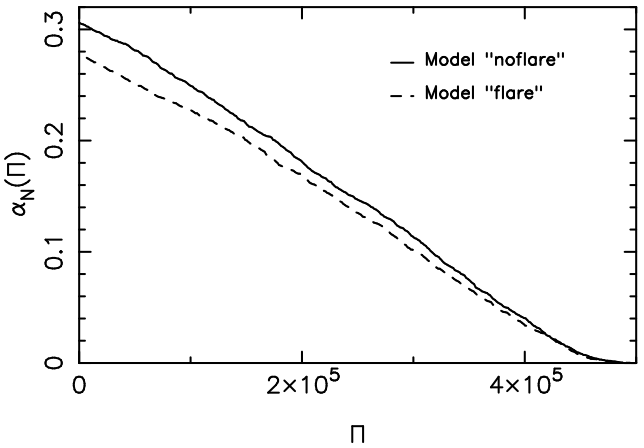

Figure 4. The p-value curves $\alpha_N(\Pi) = F^*_B\left(\sqrt{N}(\hat{D}^2 - \Pi)\right)$ for our two parametric models of the MW luminosity density distribution. Model "flare" (dashed line) is better than model "noflare" (full line) under the assumption that none of the models holds true for the data.

and determine numerically the SMDE $\hat{\vartheta}_{\mathbf{T}} = \mathrm{argmin}_{\vartheta\in\Theta}||\hat{g}^{MW} - \mathbf{T}\omega_\vartheta||^2_2$, where $||\cdot||_2$ denotes the usual L^2-norm. Finally the minimising value

$$\hat{D}^2 = ||\hat{g}^{MW} - \mathbf{T}\omega_{\hat{\vartheta}_{\mathbf{T}}}||^2$$

is computed.

To this end we use the Marquardt-Levenberg-algorithm (Press et al., (1994)) for the minimisation in a two-step process:

1. Fitting of the disc parameters: In the first step we fit the disc parameters and the bulge normalisation b, with the other bulge parameters fixed.
2. Fitting of the bulge/bar parameters: In the second step we fix the disc related parameters found in the first step (except for the normalisation parameter d) and fit the bulge/bar parameters and d.

Model	bulgemodel	discmodel	
	"BGS-bulge"	"BGS-disc"	"Kent-disc"
ρ^{noflare}	X	X	
ρ^{flare}	X		X

Table 1. Combinations of the bulge/bar and disc model components to models.

4.4 *P*-value curve analysis of our MW models

Finally we analyse the MW models with our new method. To this end we first apply the fitting algorithm (sect. 4.3) to both model "flare" and "noflare". From this we obtain the best-fitting functions $f_{\hat{\vartheta}_{\mathbf{T}}}^{\text{flare/noflare}}$, and the corresponding distances of the models from the data $\hat{D}^{2,\text{flare/noflare}}$. Here superscript flare/noflare indicates that we determine this functions for model "flare" and (separately) for model "noflare". Then we perform the bootstrap analysis (sect. 2.2) for both model "flare" and "noflare", using the corresponding best fit function and $\hat{D}^2$ for the respective model, and $B = 5000$. This yields the empirical cumulative probability distribution functions $F_B^{*,\text{flare/noflare}}$ of $X = \sqrt{N}\left(\hat{D}^{2,\text{flare/noflare}} - \Pi\right)$, and thus the functions $\alpha_N^{\text{flare/noflare}}(\Pi)$. Fig. 3 presents the resulting empirical distribution functions $F_B^{*,flare/noflare}$, and Fig. 4 the p-value curves $\alpha_N^{\text{flare/noflare}}(\Pi)$.

The last step in the analysis is to perform the graphical analysis of the p-value curves shown in Fig. 4. For every assumed distance Π between model and the "true" function f we find more evidence for model "flare" than for model "noflare" (cf. sect. 3). Therefore we are in the situation of graph 1 in fig. 2 and conclude that the p-value curve of model "flare" yields significantly more evidence for this model than that of model "noflare". Hence inclusion of flaring in the stellar disc improves the model. We can exclude that this conclusion is due to over-fitting, because the entire p-value curve performs better.

However note that the present analysis provides much more information than in [BM1], namely that there is *more statistical evidence for* "flare" as for "noflare" and *not only less evidence against* "flare" compared to "noflare". This is because the method in [BM1] is based on the assumption that the model holds (as essentially all classical goodness of fit procedures do), and therefore "only" helps to decide whether the model should be rejected given the observations. In contrast to this, for the new method proposed in this paper we assume that "the model does not hold", and estimate the probability that the distance between model and the "true" function f is smaller than any assumed distance Π. Thus variation of Π allows to find a (statistical) upper bound for the distance between model and data, providing evidence *for* the model (within the limits of the chosen distance between model and the true function ρ).

We conclude that *inclusion of flaring in the double-exponential disc improves the fit to the COBE/DIRBE L-band data.* This result is non-ambiguous, in particular since the curve of model "flare" is below the curve of model "noflare" over the entire scenario of possible distances Π in fig. 4.

Determining Π such that $\alpha_N(\Pi = D^2) \approx 0.1$ yields an estimate for the distance between the best models with and without flaring disc component, respectively, and the true density distribution of the MW. We find $\Pi \approx 3 \times 10^5$. This value can be considered as scientifically negligible because it is approximately equal to the distance $\tilde{D}^2$ between the the best fitting parametric model of [BGS] and a variant thereof in which the parameters have been changed in a random way by only $\approx 1\%$.

5 FINAL REMARKS AND CONCLUSIONS

5.1 Statistical Methodology

We have suggested a method which allows to assess the validity of a regression model at a controlled error rate. The error rate is fixed by deciding how large D^2 may at the most be while still being considerable as scientifically negligible. Furthermore, several models can be compared. This comparison is still possible if the models are parametrised by different numbers of parameters since our method is sensitive to over-fitting of the data. It is worthwhile to comment briefly on possible relationships to other approaches. As pointed out by a referee, our approach is based on weighted least squares and hence, in a model with normal heteroscedastic errors, this is the maximum likelihood estimator (MLE). Note, however, that our approach does not require the assumption of a normal error, in general.

Other approaches in the literature are based on Bayesian ideas, e.g. model averaging where the aim is to maximise the aposteriori probability of a model U_k, say, given the observations Y, i.e.

$$P(T|Y) = \sum_{k=1}^{l} P(T|U_k, Y)P(U_k|Y) \tag{5}$$

where l models are to be compared and $P(U_k|Y)$ denotes the posterior probability of the model U_k given Y. Here $T \hat{=}$ "pick the correct model" (see Hoeting et al., 1999, DiCiccio et al., 1997). This approach is conceptually similar to ours, because it is based on the idea that the decision in

favor of or against a model should be investigated under the full scenario of possible models. Bayesian model averaging aims for this by averaging, whereas we compare all P-value curves among each other. However, in addition, we are in the position to decide whether the most appropriate model by such a rule should be chosen at all. Interestingly Hoeting (p. 399) points out that such an investigation for Bayesian model averaging would be of great interest. Another difficulty in Bayesian model selection consists in the determination of priors. Observe, that our approach is based on a limit theorem, which holds for any error distribution of ε, provided $\mathrm{Var}[\varepsilon] < \infty$. It would be important to investigate more closely these relationships, however this is beyond the scope of this paper.

5.2 Flaring of the MW disc

As an example application of our method we have compared a parametric model of the MW luminosity distribution with a flaring vertical disc scale-height z_0 with a model without flaring in the disc. We find that the model with a flaring disc fits better the *COBE/DIRBE* L-band data than the model with constant z_0. We conclude that the stellar disc flares outside some inner radius R_i, which is significantly smaller than the radius of the solar orbit $R_\odot$.

Can young supergiant stars unrelated to the bulk of the stellar population, or polycyclic aromatic hydrocarbon $3.3\mu m$ or dust emission be responsible for our result? Probably not, because Alard (2000) finds flaring of the disc outwards of the solar orbit from *star count data*. It seems improbable that near the solar circle the cause of the probably same phenomenon changes. Also it is believed that the NIR luminosity probes the density of stars [BGS].

6 ACKNOWLEDGMENTS

The authors are indebted to O. Gerhard and the referee E. Feigelson for many helpful comments. N. Bissantz acknowledges support by the Swiss Science Foundation under grant number 20-56888.99, and thanks the University of Basel for their support, parts of this paper were written there. The source code can be obtained from the authors on request. Please send an email to *bissantz@math.uni-goettingen.de*

REFERENCES

Akaike H., 1974, IEEE Trans. Autom. Control AC-19, 716
Alard C., 2000, preprint (astro-ph/0007013)
Berger J. O., 1985, Statistical decision theory and Bayesian analysis. 2nd ed, Springer Series in Statistics. Springer, New York, NY
Binney J., Gerhard O., 1996, MNRAS, 279, 1005
Binney J., Gerhard O., Spergel D., 1997, MNRAS, 288, 365 [BGS]
Bissantz N., Gerhard O., 2001, MNRAS, 330, 591
Bissantz N., Munk A., 2001, A& A, 376, 735 [BM1]
Bissantz N., Englmaier P., Binney J., Gerhard O., 1997, MNRAS, 289, 651
Burnham K. P., Anderson D.A., 1998, Model selection and inference. A practical information theoretic approach. Springer, New York, NY
DiCiccio T.J., Kass R.E., Raftery A.E., Wasserman L., 1997, J. Americ. Stat. Assoc., 92, 903
DwekE., et al., 1995, ApJ, 445, 716
Efromovich S., 1999, Nonparametric curve estimation. Methods, theory, and applications. Springer Series in Statistics, New York, NY
Freudenreich H.T., 1998, ApJ, 492, 495
Gilmore G., Reid N., 1983, MNRAS, 202, 1025
Hart J., 1997, Nonparametric smoothing and lack of fit tests. Springer Series in Statistics, New York, NY
Hoeting J.A., Madigan D., Raftery A.E., Volinsky C.T., 1999, Statistical Science, 14, 382
Kent S.M., Dame T.M., Fazio G., 1991, ApJ, 378, 131
Lucy L.B., 1994, RvMA, 7, 31
Malhotra S., 1995, ApJ, 448, 138
Merrifield M.R., 1992, AJ, 103, 1552
Munk A., 1999, Validation of regression models, Habilitationsschrift, Ruhr-Universität Bochum
Munk A., Czado C., 1998, Journ. Roy. Statist. Soc. Ser. B, 60, 223
Munk A., Ruymgaart F., 1999, preprint
Press W.H., Teukolsky S.A., Vetterling W.T., et al., 1994, Numerical recipes in C (2nd ed.). Cambridge University Press, Cambridge
Schwarz G., 1978, Ann. Statist., 6, 461
Seber G.A.F., Wild C.J., 1989, Nonlinear regression. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons, Chichester
Spergel D.N., Malhotra S., Blitz L., 1995, in Minniti D., Rix H.-W., eds, Spiral Galaxies in the Near-IR. Springer, Berlin, p. 128
Stute W., Gonzáles Manteiga W., Presedo Quindimil M., 1998, Journ. Amer. Statist. Assoc., 93, 141
WeilandJ.L., et al., 1994, ApJ, 425, 81